IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

IN RE THE APPLICATION OF:	William Salkewicz	ART UNIT:	2142
APPLICATION NO.:	10/020,388	EXAMINER:	Thong H. Vu
FILED:	December 14, 2001	CONFIRMATION NO.	4901
TITLE:	DYNAMIC BINDING OF NETWORK SERVICES		

Mail Stop: Amendment
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

SUBSTITUTE PAPER

Sir,

Attached hereto is the replacement copy of the paper of Smith (with addendum) as referenced in the accompanying response to Office Action.

The paper of Smith was previously cited by the Examiner.

The paper is being replaced only because the copy previously supplied by the Office was incomplete / illegible. It is already part of the record.

Respectfully submitted,
BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP

Date: 7/5/, 2005

H H Black

H. Henry Black
Reg. No. 43,320

12400 Wilshire Boulevard
Seventh Floor
Los Angeles, California 90025-1026
(408) 720-8300

I hereby certify that this correspondence is being deposited with the United States Postal Service as first class mail with sufficient postage in an envelope addressed to the MS: Amendment, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450, on July 5, 2005.
(Date of Deposit)

Jane Wolfe
(Printed name)

Jane Wolfe — July 5, 2005
(Signature) (Date)

Visual Segmentation and the Dynamic Binding Problem: Improving the Robustness of an Artificial Neural Network Plankton Classifier (1993) (Make Corrections) (1 citation)

Graham D Smith
Centre for Intelligent Systems, University of Plymouth

bstract: A visual segmentation mechanism for a connectionist pattern recognition system sought. However, to find such a device requires the solution of the dynamic binding problem. A visual segmentation mechanism r a connectionist pattern recognition system is sought. However, to find such a device requires the solution of the dynamic inding problem. Visual segmentation could be learned by a dynamic binding network. Several putative dynamic binding echanisms are discussed but each is found to... (Update)

ited by: More
owards a Psychologically Plausible Dynamic Binding Mechanism - Smith (Correct)

ctive bibliography (related documents): More All
.5: Algorithms For Automatic Tree Ring Identification And Measurement - By Mark (Correct)
.2: On the Representation of Object Structure in Human Vision.. - Edelman, Newell (1998) (Correct)
.2: Perseverative and Semantic Influences on Visual Object Naming .. - Plaut, Shallice (1992) (Correct)

imilar documents based on text: More All
.3: Identification Of Nine Species Of Dinoflagellates, Performance .. - Contract Mas (Correct)
.3: Dinophysis - Plankton Species (Correct)
.2: The Coexistence of Plankton Species with Various - Nutrient Conditions Lattice (2004) (Correct)

ibTeX entry: (Update)

mith, G.D. (1993a). Visual segmentation and the dynamic binding problem: Improving the robustness of an artificial neural etwork plankton classifier. Technical Report, CT92-15-003. Centre for Intelligent Systems, University of Plymouth. ttp://citeseer.ist.psu.edu/smith93visual.html More

```
techreport( smith-visual,
 author = "G. Smith",
 title = "Visual segmentation and the dynamic binding problem: Improving the robustness
   of an artificial neural network plankton classifier",
 number = "CT92-15-003",
 institution = "Centre for Intelligent Systems, University
   of Plymouth",
ear = "1993",
 url = "citeseer.ist.psu.edu/smith93visual.html" }
```

itations (may not include all citations):

468 Learning internal representations by error propagation (context) - Rumelhart, Hinton et al. - 1986
13 Oscillatory responses in cat visual cortex exhibit inter--co.. (context) - Gray, Konig et al. - 1989
02 A feature integration theory of attention (context) - Treisman, Gelade - 1980
7 Tensor product variable binding and the representation of sy.. (context) - Smolensky - 1990
8 Syntactic transformations on distributed representations - Chalmers - 1990
4 Learning recursive distributed representations for holistic .. - Chrisman - 1991
1 Implications of recursive distributed representations (context) - Pollack - 1989
binding and illusory conjunctions (context) - Horn, Sagi et al. - 1991
Dynamic binding: A basis for the representation of shape by .. (context) - Hummel, Biederman - 1990
Personal correspondence (context) - -- - 1993
What connectionist models learn: Learning and representation.. (context) - MAST--CT, technical et al. - 1990
Learning to segment images using dynamic feature binding (context) - University, Zemel et al. - 1991

ocuments on the same site (http://www.cis.plym.ac.uk/cis/tech_reps.html): More
utomatic Methods for the Detection of Harmful Algae - By Culverhouse (Correct)
age Complexity Measure - Contract Mas (Correct)
ontract: MAS2--CT92--0015 - By Gavin (Correct)

CiteSeer Find: visual segement dynamic binding

Searching for **PHRASE visual segement dynamic binding**.
Restrict to: Header Title Order by: Expected citations Hubs Usage Date Try: Google (CiteSeer) Google (Web) CSB DBLP
No documents match Boolean query. Trying non-Boolean relevance query.
500 documents found. **Order: relevance to query.**

A Syntactic Theory of Dynamic Binding - Moreau (1997) (Correct) (1 citation)
A Syntactic Theory of **Dynamic Binding** To appear in the proceedings of
www.mmrg.ecs.soton.ac.uk/publications/archive/moreau1997e/moreau1997e.ps

A Lambda-Calculus for Dynamic Binding - Dami (1997) (Correct) (18 citations)
A Lambda-Calculus for **Dynamic Binding** Laurent Dami Centre Universitaire
Gen`eve 4, Switzerland dami@cui.unige.ch **Dynamic binding** is a runtime lookup operation which extracts
with flexible software assembly involve **dynamic binding**: first-class modules, mobile code,
cui.unige.ch/pub/dami/dynBind.ps.Z

Using Visualization to Foster Object-Oriented Program.. - Jerding, Stasko (1994) (Correct) (10 citations)
Using **Visualization** to Foster Object-Oriented Program
ftp.cc.gatech.edu/pub/gvu/tr/1994/94-33.ps.Z

Visual Segmentation and the Dynamic Binding Problem: Improving the .. - Smith (1993) (Correct)
Mast-Ct92-0015 Technical Report: Ct92-15-003 **Visual** Segmentation And The **Dynamic Binding** Problem
Report: Ct92-15-003 **Visual** Segmentation And The **Dynamic Binding** Problem Improving The Robustness Of An
Ct92-15-003 **Visual** Segmentation And The **Dynamic Binding** Problem Improving The Robustness Of An
www.cis.plym.ac.uk/MAST_Reports/ct92-15-003.ps.gz

Taming Message Passing: Efficient Method Look-Up for.. - Vitek, Horspool (1994) (Correct) (12 citations)
Message Passing: Efficient Method Look-Up for **Dynamically** Typed Languages Jan Vitek 1 and R. Nigel
and the message name (selector)Conceptually this **binding** of a message to its implementation proceeds as
of the object-oriented language, this **binding** can be effected at run-time, at compile-time, or
cui.unige.ch/OSG/people/jvitek/Publications/ecoop94.ps.gz

M-RPC: A Remote Procedure Call Service for Mobile Clients - Bakre, Badrinath (1995) (Correct) (5 citations)
mobile applications will also need the ability to **dynamically bind** mobile clients to local information
will also need the ability to **dynamically bind** mobile clients to local information servers. In
network with the additional features of **dynamic binding**, disconnected operation and call retries from
paul.rutgers.edu/pub/badri/mrpc.ps.Z

A Configurable Protocol Architecture for CORBA Environments - Crane, Dulay (1997) (Correct) (6 citations)
ability to select the elements of a protocol stack **dynamically** at **bind**-time depending on the properties of
the elements of a protocol stack **dynamically** at **bind**-time depending on the properties of the interface
of communication paths between its components, or **binding**. Essentially this involves the initialisation of
scorch.doc.ic.ac.uk/~jsc/research/isads-97.ps.gz

The Design of a Completely Visual Object-Oriented.. - Citrin, Doherty, Zorn (1994) (Correct) (4 citations)
To appear in **Visual** Object-Oriented Programming: Concepts and
ftp.cs.colorado.edu/pub/techreports/zorn/VOOP-VIPR.ps.Z

Cooperation Contracts - Schrefl, Kappel (1991) (Correct) (2 citations)
under grant GR 21/96106/5 ffl Overriding and **dynamic binding**: The same method may be implemented
grant GR 21/96106/5 ffl Overriding and **dynamic binding**: The same method may be implemented differently
specialization, overriding and **dynamic binding** assist in maintaining and extending designs.
linux.eecs.umich.edu/.5/groups/gasm/cooperate.ps.gz

Visualizing the World-Wide Web with the Navigational View.. - Mukherjea, Foley (1995) (Correct) (36 citations)
Visualizing the World-Wide Web with the Navigational
filtering is similar to the concept of **dynamic** queries [14]5 Hierarchization The previous

ABSTRACT

A visual segmentation mechanism for a connectionist pattern recognition system is sought. However, to find such a device requires the solution of the dynamic binding problem. Visual segmentation could be learned by a dynamic binding network. Several putative dynamic binding mechanisms are discussed but each is found to have weaknesses. Two mechanisms are being studied in greater depth so that their weaknesses may be resolved. Also a micro-world for the simulation of visual segmentation tasks is described.

1. INTRODUCTION

Simpson *et al.* (1992) have shown that an artificial neural network (ANN) can discriminate between pre-processed images of *Ceratium arcticum* (Ehrenberg) and *Ceratium longipes* (Bailey), two dinoflagellate plankton species. The input patterns on which the network was trained were outline drawings of plankton specimens taken from photomicrographs or camera lucida images. Each outline drawing was digitised and the frequency histogram of the image's power spectrum was determined by a Fast Fourier Transform (FFT). The frequency gradient of the lowest 16 frequency bins of the histogram comprised the input to the network.

However, the ANN plankton classifier is not able to classify plankton specimens contained in images where more than one specimen is present (see Figure 1) or where there are also large items of debris such as fragmants of broken plankton or air bubbles. As a result the network is trained and tested with images of single plankton specimens with no large items of clutter. This approach, therefore, fails to address an important challenge for machine vision. This challenge is to enable artificial vision systems to deal with visual images which contain more than one object. In contrast to the ANN plankton classifier, human vision segments an image into its constituent objects.

The objective of this study is to improve the robustness of the ANN plankton classifier by incorporating a mechanism which enables recognition of an object separate from the recognition of other objects contained within an image. This mechanism must be able to segment an image into its constituent objects, then generate a representation in which the information relating a single object is grouped together and kept separate from information about other objects. Psychological and connectionist theory will inform the possible mechanisms considered. In particular a general purpose segmentation mechanism is sought, not one that is only able to segment images of plankton. An exciting possibility is a mechanism which learns for itself how to segment images from a particular domain.

Firstly this report will comment on the psychology of visual segmentation, in which failures of the process are thought to yeild insights into how human brains perform segmentation. Next, a link between segmentation and connectionism's dynamic binding problem is established and some putative binding mechanisms discussed. Then the progress towards finding an appropriate segmentation mechanism is considered and the future directions of this research are identified.

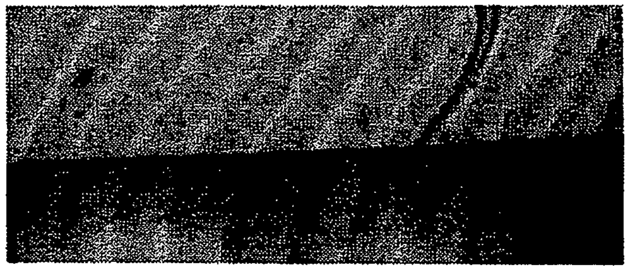

Figure 1. A photomicrograph of two Ceratium arcticum and debris.

2. VISUAL SEGMENTATION

Segmentation not only occurs in vision but is fundamental to all perceptual processes. For instance during auditory language perception sound is grouped spatially to the location of its source and is grouped temporally into phonemes and words. Treisman & Galade (1980) describe segmentation as "the process of grouping of information over the spatial extent of an object ... so that features belonging to one object are not confused with those of another object"(p 97). Unfortunately this definition only considers those objects which extend spatially. It ignores the possibility of objects, such as sounds, which extend over time. Therefore, Treisman & Galade have only defined visio–spatial segmentation. Also it is important to distinguish between two possible uses of the term *object*. Object_1 refers to an object as it exists in the physical world, whereas object_2 refers to the mental representation a perceiver has of an object_1. Treisman & Galade's definition uses *object* in the sense of object_1. In this report the term *object* should be taken to mean object_1 and the term *perceptual object* should be understood as object_2. Also for the purposes of this report the components of a representational object shall be referred to as elements or representational elements.

Perceptual processes are so reliant on segmentation that it is difficult for us to imagine what unsegmented experiences would be like. Only under very unusual circumstances do errors of segmentation occur. One example which gives an impression of what unsegmented perception might be like occurs when one looks at a television screen from close up. At a distance of a few centimetres the picture cannot be discerned. All that can be seen is a mosaic of red,blue and green dots. At this distance from the television our visual knowledge is insufficient to group these coloured elements together to represent the objects displayed on screen. The reason why, is that this is a visual experience very different from that which we normally encounter. Also as one moves back from the television screen there comes a point were one is able to group together the dots of coloured light and segment the image into objects. This point is where the image becomes enough like the normal visual world for our visual knowledge to enable segmentation.

Segmentation and object recognition are intimately related, in fact it appears that in order to work both processes require information from each other. It is possible that the processes of segmentation and object recognition occur in parallel. Obviously, an object can only be recognised once it has been isolated from an image. But also accurate segmentation relies on knowledge about the internal structure of objects. The *too close to the TV* example, mentioned earlier, demonstrates that in situations were a person's visual knowledge is limited by lack of experience they are unable to segment effectively. However,we are able

toscopically presented with a visual array of coloured letters and asked to respond if a certain letter and colour combination is present in the stimulus. It is found that subjects will frequently respond when the colour and the letter are both present in the array but are separate. Theories of visual segmentation must account for the occurrence of false conjunctions.

The search for a segmentation mechanism which can be instantiated in a connectionist network is a facet of the dynamic binding problem. There now follows a discussion of dynamic binding and the types of connectionist representations which have been proposed as solutions to the problem.

3. THE DYNAMIC BINDING PROBLEM

Dynamic binding is the representation of objects by the temporary conjunction of two or more representational elements. It is required by any cognitive process that exhibits systematicity and compositionality (Shastri & Aajanagaade, in press), processes which include the production and comprehension of language, logical inference and visual segmentation. For example, Treisman and Galade's definition of visual segmentation includes the term "the grouping of information". This term is equivalent to binding. But connectionist theory will be unable to account for systematic and compositional cognitive processes unless a mechanism which can perform dynamic binding in ANNs can be found. However, none of the binding mechanisms devised to date are fully satisfactory. There shortly follows a discussion of these putative binding mechanisms which will highlight each mechanism's strengths and weaknesses. These binding mechanisms fall into four categories: associationist representations, enumerated representations, phase synchrony in oscillatory networks and recursive distributed representations (RDR).

3. 1. Associationist Representations

Associationist representations are the simplest of connectionist representations. They identify those features which are present but not how the features are grouped into objects. Therefore, they are unable to simultaneously represent more than one object. It is this weakness of associationist representations which brought attention to the dynamic binding problem for connectionism. To see how this type of representation fails to describe two objects at once the following test will be applied. The test is to represent the simplest possible two object description. Each object is the conjunction of two representational elements: a colour and a geometric shape. There are two possible colours, i.e. red or blue, and two possible shapes, i.e. square or triangle. If a putative dynamic binding mechanism is able to represent a red triangle and a blue square then it is worthy of consideration. An associationist representation of *red triangle and blue square* would have active units standing for the elements *red*, *blue*, *square*, and *triangle*. However there is no difference between the associationist representation for *red triangle and blue square* and the representation for *blue triangle and red square*. Therefore, associationist representations are unable to perform binding.

3. 2. Enumerated Representations

Enumeration is the most widely used binding technique. It is a representation in which there are units standing for each of the possible conjunctions of representational elements. For example, an enumerated representation which allows the representation of *red square and blue triangle* would have a unit for each pairing of colour and shape i.e. *red square*, *blue square*, *red triangle* and *red square*. This is a localist enu-

Another major weakness of enumerated representation is that the representational elements are unbound. For instance in an enumerated representation of *red triangle and red square* there is nothing to indicate that the *red* in *red square* is the same as the *red* in *red triangle*. This is the equivalence problem of enumerated representations. Enumeration does not dynamically bind representational elements together at all. These elements are not dynamically bound they are permanently linked.

3. 3. Phase Synchrony in Oscillatory Networks

In this technique representation of binding is achieved by allocating bound elements, to the same phase in an oscillatory network. Oscillatory networks have units whose activation fluctuates over time. Binding is represented using the phase of the waveform of this oscillation of activation. The elements of a representational object all have the same phase.

Support for phase synchrony binding comes from neuropsychological evidence which shows that the activity of non–adjacent hypercolumns in a cat's visual cortex will become phase locked when perceiving different parts of the same moving object (Gray *et al.*, 1989). Also it is claimed that limits on the number of possible discrete phases an oscillating system can maintain may explain some constraints on the cognitive abilities of humans. It is argued that accidental synchrony between units can account for false conjunctions (Hummel & Biederman, 1990) and for the 7+2 constraint on Short–term memory (Shastri & Ajjanagadde, in press).

Most oscillatory networks have complicated mechanisms for producing the oscillation of activation and for establishing phase synchrony. For instance, Hummel & Biederman's (1990) network co–ordinates the phases of units by activating fast enabling links. These are links which operate on a time–scale several times faster than that of ordinary connections. Many oscillatory networks have activation waveforms that are convoluted and even chaotic at times(e.g. Horn, Sagi & Usher, 1991). The complexity of these networks and their behaviour are barriers to our understanding of phase synchrony as a binding mechanism.

In contrast to these complicated networks, the oscillatory activation of Mozer, Zemel & Behrmann's (1992) MAGIC is implicit. In this network, activation is represented by complex numbers in polar form. The amplitude of activation represents confidence in the presence of an element. The phase of activation represents binding between elements. MAGIC has been trained to segment images of a micro–world comprised of simple geometric shapes. Representations of these shapes are constructed using four types of feature: lines at 0°, 45°, 90° and 135° to the vertical. There is a unit for each possible conjunction of feature and location which are grouped as feature maps. When an input pattern is first presented, the phase of each unit is random. The network then relaxes to a stable pattern of activation in which elements from the same object have the same phase. Also the network has a learning rule which is able to find a set of complex valued weights that will perform this segmentation. But despite the success of MAGIC, phase synchrony is a poor dynamic binding mechanism because it is unable to represent shared features.

Existing oscillatory networks do not allow an element to be part of two objects at the same time because the unit representing such an element would need to have two phases, which is not possible. For instance to represent *red circle and red square*, the element for *red* must be bound to element for *square* whose unit has one phase and be bound to the element for *circle* whose unit has a different phase. Therefore the *red* unit requires two phases. A new kind of oscillatory network would be needed to overcome this problem of representing shared elements.

3. 4. Recursive Distributed Representations

The failure of all of the above–mentioned putative dynamic binding mechanisms is possibly founded on their use of a restricted compositional method, *concatenative compositionality* (van Gelder, 1990). This form of compositiality is the construction of structured representations from a set of tokens (i.e. representational elements) by linking or ordering them so that these tokens are unchanged. For example in propositional logic the statement *(p&q)* retains the tokens *p* and *q*. However, van Gelder also argues that preservation of tokens is not necessary for compositionality. There is also *functional compositionality* in which complex expressions do not contain tokens of their constituents, although these constituents are retrievable.

Pollack's (1989) Recursive Auto–Associative Memory is able to form representations which appear to possess functional compositionality. This type of representation is called a recursive distributed representation (RDR). RDRs, as their name suggests, are produced in a distributed network by a recursive process. Over several serial steps a global representation of a group of elements is produced by adding a further element to the global representation produced by the last step. An RAAM would have no trouble in representing *red triangle and blue square*. Up until now RDR had not been proposed as dynamic binding mechanism.

However there is a weakness of the RDR approach to visual segmentation. It is very unlikely that visual segmentation is a serial process but RDRs are composed and decomposed serially. It takes us only a fraction of a second to make sense of a new visual image. The speed of this process strongly suggests that the brain computes visual segmentation in parallel. However, the representation developed by RAAMs on their hidden units is able to encode binding of elements despite the serial process by which the hidden activation representation were formed. Therefore, understanding how objects are represented within RDRs may help us to solve the dynamic binding problem.

4. PROGRESS

The progress made on this project has been divided between three areas. The first of these areas has been to design a micro–world which approximates to real–world visual segmentation. Another area of progress has been the development of Polarnet: an complex domain oscillatory network. The final area of progress was a preliminary study of a RAAM's performance when trained with patterns taken from the micro–world. These areas of progress and their future development will be discussed in turn.

4. 1. Micro–world

The segmentation of real–world images is a complex and computationally intensive process which is far beyond the capability of networks comprised of only a few tens of units. Therefore, a micro–world has been developed images from which are much easier to segment than real–world images. It is intended that binding networks be trained and tested with micro–world images to assess their viability. A micro–world is a simulated universe which is defined by a few simple rules. These rules, which can be thought of as the micro–world's *laws of nature*, should reflect those characteristics of the real–world which may influ-

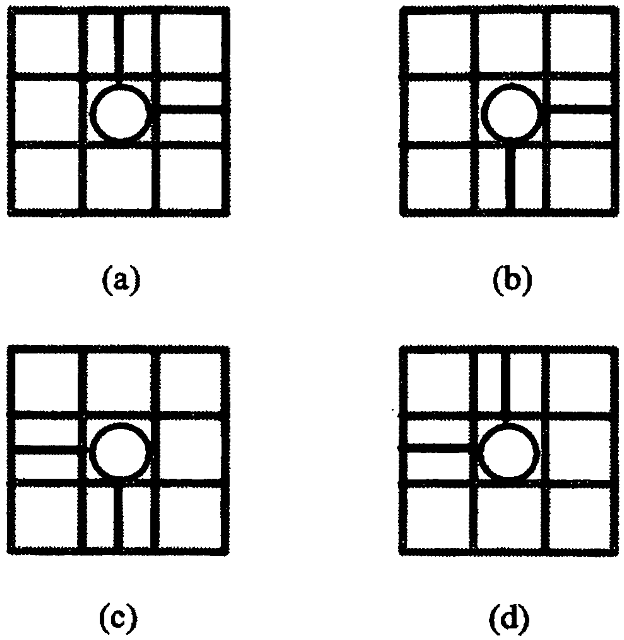

Figure 2. Four Micro–world Plankton at Different Rotations

4. 1. 2. Boundary Problems

A micro–world of a visual domain must define how objects are represented when some of the objects features are located over the micro–world's horizons. One possible interpretation is to disallow those objects which cross a micro–world horizon in this way. However this interpretation causes problems. As a result certain feature – location conjunctions are more common than others. This situation is definitely not true of human vision. Also some of these conjunctions are more frequently associated with a particular species of micro–world plankton which becomes the basis by which plankton species are categorized. This is quite unlike real–world categorisation. A more realistic categorisation task is based on the recognition of a spatial arrangement of features within an object. This sort of categorisation will occur in the absence of simpler cues to recognition. To ensure the even distribution of feature to location conjunctions those features of objects which disappear over the microworld horizon will reappear over the opposite horizon (See Figure 3). This interpretation permits 36 different plankton exemplars.

(a) (b)

Figure 3. Micro–world plankton whose features cross over the domain's horizon.

4. 1. 3. Images of Multiple Plankton Exemplars

Plankton–world allows the simultaneous representation of more than one plankton in an image. There are 630 possible two object images and over 23,000 three object images. Even in this simple micro–world the number of combinations of objects are large. When more than one object is in an image the possibility of them overlapping occurs. Figure 4 shows pairs of plankton in different degrees of overlap permitted by plankton–world.

4. 2. Polarnet

Polarnet has been inspired by MAGIC (Mozer *et al.*, 1991) but there are several key differences between them which shall be highlighted in the following description of Polarnet.

4. 2. 1. Architecture

A crucial difference between Polarnet and MAGIC is that Polarnet is a multilayer feed–forward network whereas MAGIC is a recurrent network. Therefore it will be simpler to analyse Polarnet's performance than MAGIC's. Recurrent networks develop hidden layer representations which defy interpretation. It is envisaged that existing techniques for understanding the representations developed by real–valued feed–forward networks (e.g. Hanson & Burr, 1990) will be modified for the analysis of Polarnet.

4. 2. 2. Learning

Polarnet, like MAGIC is trained through the back–propagation of error. However, the error measure used in Polarnet is unlike the error measure employed in MAGIC. Mozer *et al.* use an error measure which appears to be based on the Hamming distance between the network's actual and target outputs. Polarnet's error measure which is shown in Equation 1 is the generalisation to the complex domain of the error measure which underlies the delta or Widrow–Hoff learning rule.

$$E = 1/2 \sum_j (r_j^2 + s_j^2 - 2\, r_j s_j \cos(\Phi_j - \theta_j)) \qquad (1)$$

where E is error, r_i is the actual amplitude of unit i, s_i is the target amplitude of unit i, Φ_i is the target phase of unit i and θ_i is the actual phase of unit i. The learning rule is derived by partial differentiation of E with respect to W_{ij}, where W_{ij} is the complex valued weight of the connection between unit i of the hidden layer and unit j of the output layer. The learning rule is the negative product of this partial differential and a real valued constant which controls the rate of learning. Rumelhart, Hinton and Williams (1986) have shown that the error between the output and target activation vectors can be propagated back through the network to amend the weights of connections in all the layers. However the derivation of the backpropagation learning rule must take account of a network's activation rule.

(a) (b) (c)

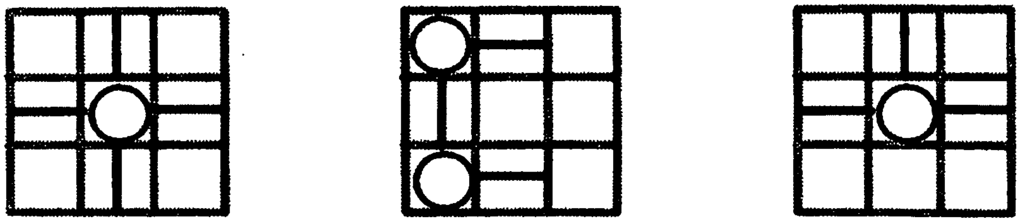

(d) (e) (f)

Figure 4. Examples of micro–world images in which there are different types of overlap between two Plankton; (a) is of non–overlapping objects, (b) and (c) are of two objects with different features at the same location, (d) and (e) are of two objects with the same feature at the same location, and (f) is of two objects with two of the same features at the same locations.

4. 2. 3. Activation

Polarnet's net input function is the dot product of the activation and weight vectors generalised to the complex domain, which is the same as MAGIC's input function. However, no satisfactory output function has been found for Polarnet as yet. The network requires an output function which maps the net input into a unit to a complex number in the range $0 < r_i < 1$ and $-\pi < \theta_j \leq \pi$. Values for r_j and θ_j, which are outside of this range cannot be defined in polar–form complex numbers. The output function should also be non–linear and allow partial differentiation with respect to W_{ij}. Meeting these conditions will allow the network to use a back propagation learning rule. An early version of Polarnet employed the logistic function shown in Equation 2 which proved to be unsatisfactory.

$$r_i = (1-e^{-|n_i|})^{-1} \qquad (2)$$

where r_i is the amplitude of unit i and n_i is the net input to unit i. In Polarnet values of $|n_j|$ are always positive and this function maps positive values to the range $1/2 \leq r_i < 1$. Consequently the network fails to learn the input–output mappings with which it is trained. A promising alternative to the logistic function is shown in Equation 3 (Georgiou, 1993).

$$f(n_j) = n(1+|n|)^{-1} \qquad (3)$$

Polarnet has yet to be implemented with this activation function.

4. 2. 4. Future Work

Once a suitable activation function has been found simulations of Polarnet will be undertaken to study the network's ability to segment images taken from Plankton–world. The networks performance will be

Alongside the development of Polarnet, investigation of a recursive distributed representation approach to visual segmentation of Plankton–world images has been undertaken. The input pattern representing an unsegmented image is divided into several sub–patterns which are presented serially to a RAAM. Each sub–pattern stands for the conjunction between a feature and its location. The sequential order of sub–patterns denotes the object to which the feature belongs.

For a RAAM to map a random sequence of features to a segmented sequence of features, a transformation of the network's hidden representations is required. Several researchers have simulated RAAM networks in which such transformations are made (Chalmers, 1990; Chrisman, 1991). The transformation process can be facilitated as follows. A RAAM is trained to auto–associate input patterns taken from the micro–world. These patterns are presented as a randomly ordered series of sub–patterns from which the network recursively generates a global representation of the sequence. After training the network's hidden representation for each input pattern is found. A second RAAM network is trained to auto–associate the set of target output patterns in which the features are sequentially grouped into objects. This network's hidden layer representations of the target patterns are found after training. A third network performs the transformation in which the first RAAM's representations of the input patterns are mapped to the second RAAM's representations of the segmented output. A single network can be constructed from some of the units and weights of the original networks which should map unsegmented images to segmented images. Chrisman (1991) has devised a variant of this technique which can be implimented in a single network, which he calls a dual–ported RAAM. Simulations of a dual–ported RAAM will be undertaken for comparison with Polarnet.

5. CONCLUSIONS

Segmentation and dynamic binding are seen to be intimately related when perception is studied from a connectionist viewpoint. Whilst no psychologically plausible nor entirely adequate segmentation mechanism is suggested by connectionist work on the dynamic binding problem, two approaches, phase synchrony in oscillatory networks and recursive distributed representations, do offer some promise. Investigations into the viability of these two binding mechanisms will continued.

6. REFERENCES

Chalmers, D.J. (1990). Syntactic transformations on distributed representations. *Connection Science*, 1&2, 53–62.

Chrisman, L. (1991). Learning recursive distributed representations for holistic computation. *Connection Science*, 3, 345–365.

van Gelder, T. (1990). Compositionality: A connectionist variation on a classical theme. *Cognitive Science*, 14, 355–384.

Georgiou, G.M. (1993). *Personal correspondence.*

Gray, C.M., Konig,P., Engel,A.K. &Singer,W. (1989). Oscillatory responses in cat visual cortex exhibit inter–columnar synchronization which reflects global stimulus properties. *Nature*, 338, 334–337.

Mozer, M.C., Zemel, R.S. & Behrmann, M. (1991). Learning to segment images using dynamic feature binding. *Technical report*, CU–CS–540–91. University of Colorado, Boulder.

Rumelhart, D.E., Hinton, G.E. & Williams, R. (1986). Learning internal representations by error propagation. *Nature*, 323, 533–536.

Shastri, L. & Ajjanagadde, V.(in press). From simple associations to systematic rules: A connectionist representation of rules, variables, and dynamic bindings using temporal synchrony. submitted for publication in *Behavioral and Brain Sciences*.

Simpson, R., Williams, R., Ellis, R. & Culverhouse, P.F. (1992). Biological pattern recognition by neural networks. *Marine Ecology Progress Series*, 79, 303–308.

Smolensky, P. (1990). Tensor product variable binding and the representation of symbolic structures in connectionist systems. *Artificial Intelligence*, 46, 159–216.

Treisman, A. & Gelade, G. (1980). A feature integration theory of attention. *Cognitive Psychology*, 12 97–136.

Carlos Medina
<cmedina@exponent.com>

01/25/2005 05:22 PM

To: "'Jane_Wolfe@bstz.com'" <Jane_Wolfe@bstz.com>
cc:
Subject: FW: Interpretation of paper

Jane, it appears that neither the school nor the author have copies of that paper any more. In this email, the author tries to fill in what was missing on the portions of the article that got cut off.

-----Original Message-----
From: Phil Culverhouse [mailto:P.Culverhouse@plymouth.ac.uk]
Sent: Tuesday, January 25, 2005 2:23 AM
To: Carlos Medina
Subject: Interpretation of paper

Dear Carlos,

This is the best we can do. I hope it clarifies your understanding. Are you working on a related topic??

Kind regards,

Phil

-----Original Message-----
From: Smith Graham D [mailto:Graham.D.Smith@northampton.Ac.Uk]
Sent: 25 January 2005 10:16
To: Phil Culverhouse
Subject: RE: Hello!

Phil

My memory is good but not good enough to recall all the missing details. It looks as though 6 or so lines of text have been lost from the top of each page. Much of the lost text on the early pages (e.g., the missing part of Section two which described the phenomenon of illusory conjunctions) is not needed to understand the real contribution of the paper; Polarnet (and the domain I suppose). So to sum up, the paper identifies phase synchrony in polar-form MLPs and recursive distributed representations like those used by Pollack's RAAM as potential dynamic binding representations.

The missing part of Section 3.2 appears to identify problems with enumerated representations of dynamic binding domains. What appears to be missing is discussion of the scaling problem; i.e., enumerated representations of any real world domain require unfeasibly large number of units. The equivalence problem, not previous mentioned in the literature, but related to Fodor & Pylyshn's compositionality, is described without any apparent omissions.

I recall that Figure 3 was of examples of "micro-plankton" whose "arms" cross the domain horizon and therefore appear on the otherside of the grid. This "unrealistic" arrangement was necessary to ensure an even distribution of features to locations and thereby avoid the network solving the learning problem using simple feature relationships. The number of plankton exemplars (i.e, 36) is given by the number of locations for the "micro-plankton" body (i.e., 9) multiplied by the number of possible orientations of the "micro-plankton" (i.e., 4).

I don't believe that Polarnet's learning rule was stated in the paper (there are no missing equations). I describe the polar valued version of the logistic function as being an unsatisfactory (for reasons I do not recall) activation output function. I offer an alternative function but do not

present a derivation of the learning rule from it. I have been unable to find Polarnet's learning rule in my notebooks. However, I recall that the learning rule was derived from the activation rule and error function following the strategy used by Rumelhart and McLelland as decribed in the PDP bible and Rumelhart, Hinton and Smith (1986). Also, anyone interested in polar-form NNs should look at G.M. Georgiou's work.

The paper described work in progress however subsequently my simulations focused on the RAAM network and encoders rather than Polarnet which was side-lined. My reason for this change of tack was that I realized that there is no mapping that a polar-valued activation network can perform that a normal back-prop net cannot perform. In fact, expressing the activation and learning rules in polar co-ordinates is merely a redescription of the Cartesian form of the rules. In other words if Polarnet can learn the input-output mapping then so can the equivalent backprop MLP. Better then to perform simulations on the well understood MLP than the largely unknown Polarnet.

I hope this helps.

Regards

Graham

Dr Graham D. Smith
Division of Psychology
University College Northampton
Boughton Green Road
NORTHAMPTON
NN2 7AL

Tel: +44 (0) 1604 735500 Ext 2393
E-mail: graham.d.smith@northampton.ac.uk